arXiv:1609.08562v1 [quant-ph] 27 Sep 2016

Security of two-state and four-state practical quantum bit-commitment protocols

Ricardo Loura[1,2], Dušan Arsenović[3], Nikola Paunković[1,2]*,
Duška B. Popović[3] and Slobodan Prvanović[3]

[1] Instituto de Telecomunicações, Av. Rovisco Pais 1049-001, Lisboa, Portugal
[2] Departamento de Matemática, Instituto Superior Técnico, Universidade de Lisboa, Av. Rovisco Pais 1049-001, Lisboa, Portugal
[3] Institute of Physics, University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia

Abstract

We study cheating strategies against a practical four-state quantum bit-commitment protocol [4] and its two-state variant [10] when the underlying quantum channels are noisy and the cheating party is constrained to using single-qubit measurements only. We show that simply inferring the transmitted photons' states by using the Breidbart basis, optimal for ambiguous (minimum-error) state discrimination, does not directly produce an optimal cheating strategy for this bit-commitment protocol. We introduce a new strategy, based on certain post-measurement processes, and show it to have better chances at cheating than the direct approach. We also study to what extent sending forged geographical coordinates helps a dishonest party in breaking the binding security requirement. Finally, we investigate the impact of imperfect single-photon sources in the protocols. Our study shows that, in terms of the resources used, the four-state protocol is advantageous over the two-state version. The analysis performed can be straightforwardly generalised to any finite-qubit measurement, with the same qualitative results.

1 Introduction

With the accelerating development of quantum information theory, new cryptographic protocols and primitives arise every day alongside new engineering techniques to turn them into reality. Naturally, much to the resemblance of classical cryptography and its implementations, new challenges and new difficulties appear once one moves from theory to practice. Quantifying the impact of these difficulties and studying their solutions is thus of utmost importance if quantum cryptography is to become a widespread reality. As a step towards this goal, we here study the feasibility of several different attacks on a practical two-state quantum bit-commitment (BC) protocol [10, 1, 2] and its original four-state variant [4] under a noisy environment. By practical, we mean that cheating is subjected to current technological constraints: the lack of long-term quantum memories and, in the case of the currently predominant optical realisations, the non-existence of photon non-demolition measurements. Thus, an agent (in our case Alice) is forced to perform measurements as soon as (s)he receives the qubits encoded in, say, photon polarisation. Various practical cryptographic protocols based on bounded and/or noisy memories were previously studied [16, 13, 7, 12], in order to avoid the Lo-Chau [9] and Mayers [11] no-go theorem regarding non-relativistic implementations of quantum bit-commitment schemes. (see also a recent proposal of a computationally secure oblivious transfer protocol [15]).

*npaunkov@math.tecnico.ulisboa.pt

We start by reviewing the protocol's steps, including the careful definition of the verification procedure, in Section 2. In Section 3 we discuss the so-called *cheating observable*, given by the Breidbart basis, and show that the naïve approach, studied in [10], can be perfected. In Section 4, we analyse the so-called *faking-distance attack*, based on providing wrong geographical information to the verifier. Lastly, Section 5 is dedicated to the so-called *beam-splitter attack*, which takes advantage of the fact that no perfect single-photon sources exist, much to the resemblance of BB84's photon-number splitting attack. We compare the enhanced attack with its counterpart in the four-state variant.

2 The protocol

Commitment protocols, initially proposed in 1981 by Manuel Blum [3] and simultaneously by Shamir, Rivest and Adleman [14], are nowadays important cryptographic primitives that serve as building blocks for several cryptographic and computational protocols, such as zero-knowledge proofs [5] and multi-party computation [8]. Simply put, a commitment scheme is a sequence of steps between a prover (Alice) and a verifier (Bob) that allows them to essentially roll a dice in a fair way over a communication channel. In the first step (the *commitment phase*), Alice makes a choice out of a set of possible choices ($1, \ldots, n$ in the case of an actual n-sided dice). In the second step (the *revealing phase*), Bob asks Alice to reveal her choice and to *prove* that it was indeed her initial choice. In practice, these two steps aren't immediately done after one another, but are instead crucial parts of a larger protocol. Naturally, the scheme must guarantee that between both phases Alice cannot change her mind (called the *binding* requirement), and that Bob cannot learn Alice's commitment choice before she reveals it (called the *concealing* requirement). For a more detailed description and examples of commitment schemes, see [10]. If Alice's choice is bound to be either 0 or 1, then we say the commitment scheme is a bit-commitment scheme. Furthermore, if the protocol requires quantum communication of any kind, we say that it is a quantum commitment scheme.

In this work we will be interested in a particular quantum bit commitment scheme initially presented in [10], building on the earlier works of Danan and Vaidman in [4]. We begin with a brief review of the protocol's definition. A detailed version of what follows, along with a theoretical security analysis, can be found in [10] and in [2]. Note, however, that in the present paper we use a slightly different notation. In fact, the works [10] and [2] focus exclusively on the 2-state protocol, whereas we here aim at studying both the 2-state and 4-state protocol variants. Since the original notations of each of these are incompatible, we opt here for the standard quantum computation convention, recalled below.

The two-state protocol goes as follows (the analogous four-state protocol is presented in Section 3.2):

1. Alice and Bob agree on a computational basis $\{|0\rangle, |1\rangle\}$, and, as usual, on the vectors $|+\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$ and $|-\rangle = 1/\sqrt{2}(|0\rangle - |1\rangle)$.

2. Bob randomly chooses one of the two states $|0\rangle$ or $|+\rangle$, then prepares and sends a particle in that state to Alice. He repeats this step a number M of times of his choice.

3. Alice measures either $\hat{C}_0$, for committing to 0, on all the particles she receives, or she measures $\hat{C}_1$, for committing to 1, on all the particles she receives, where $\hat{C}_0$ and $\hat{C}_1$ are the observables

defined by

$$\begin{aligned}\hat{C}_0 &= 0 \cdot |0\rangle\langle 0| + 1 \cdot |1\rangle\langle 1| \\ \hat{C}_1 &= 0 \cdot |+\rangle\langle +| + 1 \cdot |-\rangle\langle -|.\end{aligned} \tag{1}$$

4. Whenever Alice wants to reveal her commitment, she sends her commitment choice and the results of her measurements back to Bob.

5. Bob either accepts or discards Alice's commitment, based on some statistical criterion of his choosing.

Since the states sent by Bob are non-orthogonal, the outcomes of Alice's measurements follow different probability distributions, according to whether $\hat{C}_0$ or $\hat{C}_1$ was the chosen observable. Indeed, if we denote by $p_i(j|k)$ the probability that Alice measures the value j with the observable $\hat{C}_i$ whenever Bob sends a particle in state $|k\rangle$, we have

$$\begin{aligned}p_0(0|0) &= 1 \\ p_0(1|0) &= 0 \\ p_0(0|+) &= \frac{1}{2} \\ p_0(1|+) &= \frac{1}{2},\end{aligned} \tag{2}$$

and, similarly,

$$\begin{aligned}p_1(0|0) &= \frac{1}{2} \\ p_1(1|0) &= \frac{1}{2} \\ p_1(0|+) &= 0 \\ p_1(1|+) &= 1.\end{aligned} \tag{3}$$

Choosing either observable $\hat{C}_0$ or observable $\hat{C}_1$ has thus a clear impact on Alice's set of measurement outcomes, or in other words, her sets of statistical data, her statistical signature. Consequently, once Alice reveals her results, Bob can know whether she used observable $\hat{C}_0$ or observable $\hat{C}_1$ (see below). Evidently, measuring any other observable, or different observables on each photon, will destroy the statistical signature. By exactly how much, and whether or not Alice can take any kind of advantage out of it, is the topic of the next section.

Unfortunately, things are not as simple in any real-world implementation of the protocol, where any emission, transmission and measurement processes are inevitably subject to a certain amount of white noise. To account for such imperfections during the course of the protocol, we model the white noise as a depolarizing channel, or r-noise channel, $\mathcal{E}_d$, such that with probability r the state becomes totally mixed, whilst with probability $(1-r)$ it remains unharmed. It can thus be represented as a super-operator acting on a state ρ by

$$\mathcal{E}_d(\rho) = (1-r)\rho + r\frac{I}{2}. \tag{4}$$

As expected, such a channel affects the sets of probabilities (2) and (3), which now become functions of r. We shall abuse notation and refer to the new set of probabilities by the same

symbols, as the former won't be of use henceforth. We have

$$\begin{aligned} p_0(0|0) &= 1 - \frac{r}{2} \\ p_0(1|0) &= \frac{r}{2} \\ p_0(0|+) &= \frac{1}{2} \\ p_0(1|+) &= \frac{1}{2}, \end{aligned} \tag{5}$$

and

$$\begin{aligned} p_1(0|0) &= \frac{1}{2} \\ p_1(1|0) &= \frac{1}{2} \\ p_1(0|+) &= \frac{r}{2} \\ p_1(1|+) &= 1 - \frac{r}{2}. \end{aligned} \tag{6}$$

Finally, we need to formulate a systematic way for Bob to accept or discard Alice's commitment. Once Alice reveals her commitment and her measurement results, Bob's goal is to determine whether her statistics follow the probability distribution defined by (5), the one defined by (6), or none of these two. This problem, of deciding whether a sample comes from a given distribution or not, is an extremely well-known problem in Statistics, and can be solved through a variety of methods, jointly called goodness-of-fit tests, the most famous of which is likely to be the chi-squared test. We will here, however, take advantage of the simplicity of our particular case and instead use a far more straightforward test, called the binomial test.

To do so, let's assume that Bob sends N_0 particles in state $|0\rangle$. Without loss of generality, suppose Alice measures the $\hat{C}_0$ observable on all of them, i.e., she commits to the value 0. Then the probability $P(n_0\ 0\text{'s})$ that she obtains n_0 times the value 0 after measuring all the particles is given, by definition, by a binomial distribution

$$P(n_0\ 0\text{'s}) = \binom{N_0}{n_0} p_0^{n_0}(0|0) p_0^{N_0 - n_0}(1|0). \tag{7}$$

This binomial distribution has mean value μ_0 and variance σ_0^2 defined by

$$\begin{aligned} \mu_0 &= N_0 p_0(0|0) \\ \sigma_0^2 &= N_0 p_0(0|0)(1 - p_0(0|0)). \end{aligned} \tag{8}$$

The binomial test then goes as follows: Bob proceeds if Alice's statistics satisfy

$$n_0 \in [\mu_0 - 3\sigma_0, \mu_0 + 3\sigma_0]. \tag{9}$$

Otherwise, he immediately aborts the protocol. The above condition ensures that if Alice is committing to 0, then she has $\sim 99.7\%$ chance of getting her first set of statistics accepted by Bob. This value is simply

$$0.997 \sim \sum_{k=\lceil \mu_0 - 3\sigma_0 \rceil}^{\lfloor \mu_0 + 3\sigma_0 \rfloor} \binom{N_0}{k} p_0^k(0|0) p_0^{N_0 - k}(1|0). \tag{10}$$

Analogously, Bob tests Alice's statistics whenever state $|+\rangle$ was sent, and proceeds if and only if

$$n_+ \in [\mu_+ - 3\sigma_+, \mu_+ + 3\sigma_+]. \tag{11}$$

Otherwise, he immediately aborts. The probability that Alice's commitment is accepted by Bob, i.e. the probability that Alice passes both tests, is then simply the product of the separate probabilities. Note, of course, that one could choose a standard deviation factor other than 3. The smaller it is, the more restrictive the criterion becomes. For simplicity, throughout the paper we set that Bob sends equal number of particles in each state, i.e., $N_0 = N_+ = N$ (and $N_{0/1} = N_{+/-}$ in the four-state case), the assumption secured by the law of large numbers.

The test is, *mutatis mutandis*, the same if Alice were to commit to 1 instead. However, there is still a loose end: there is a nonzero probability that an Alice trying to commit to 1 produces statistics that satisfy both (9) and (11), thus convincing Bob she committed to 0, and making the test nonviable. This probability decreases exponentially with the number of measurements. Heuristically, a number of total measurements of $M = 100$ (about fifty 0's measured, and fifty 1's measured, $N = N_0 = N_+ = 50$) is enough to guarantee viability of the protocol for any reasonable amount of noise, as the plot below shows.

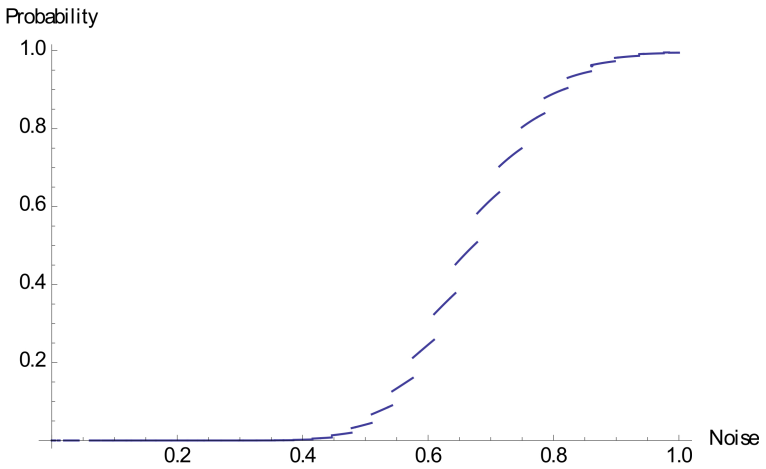

Figure 1: Probability, as a function of white noise, of an honest Alice committing to 1, with $M = 100$ successful measurements, passing Bob's test of commitment to 0.

3 Optimal cheating observable

The main goal of this section will be to answer a simple question: can Alice cheat? In other words, can she devise a strategy so that she's able to convince Bob that she made a choice when in reality she didn't? Naturally, Bob could also try to break the protocol. However, he has a much tougher job. For Bob to break the protocol, he would need to know what Alice's choice of measurement was before Alice making it public. Since he has no access to Alice's lab, doing so is impossible, even if using entangled pairs, due to non-signalling. Alice's side of things seem much more promising. For Alice to cheat, she has to be able to produce statistics that follow either distribution (5) or distribution (6) upon request. Since the only thing she can do is perform measurements on the incoming particles, she needs to find the measurement that best allows her to mimic *both* those distributions, eventually after some post-processing.

Note, of course, that, in an ideal setting, Alice could simply store the incoming particles in a quantum memory and measure them only when announcing her commitment. However, stable

long-term quantum memories are still a long way from becoming an everyday item, and we thus assume Alice is forced to perform her measurements on the incoming particles as soon as they arrive. For a more complete discussion, including the case of noisy memories, please refer to [10].

3.1 The two-state protocol

The first attempt at creating a cheating strategy is, naturally, to try and guess the state of every photon sent by Bob. Since there is no way of perfectly and systematically distinguishing non-orthogonal states, the best Alice can do, provided she is constrained to performing single-qubit measurements only, is to measure an observable "in-between" $\hat{C}_0$ and $\hat{C}_1$. Finding the basis that best discriminates states $|0\rangle$ and $|+\rangle$ is a common and well-studied problem in quantum mechanics. Such a basis is called a Breidbart basis. Specifically, let $|\tilde{0}\rangle$ (resp. $|\tilde{1}\rangle$) denote the vector $|0\rangle$ (resp. $|+\rangle$) rotated by $-\pi/8$ (resp. $\pi/8$), so as to form the aforementioned Breidbart basis, and define a new observable, the *cheating observable* $\hat{C}_{ch}$, by

$$\hat{C}_{ch} = 0 \cdot |\tilde{0}\rangle\langle\tilde{0}| + 1 \cdot |\tilde{1}\rangle\langle\tilde{1}|, \tag{12}$$

where

$$\begin{aligned} |\tilde{0}\rangle &= \cos\frac{\pi}{8}|0\rangle - \sin\frac{\pi}{8}|1\rangle \\ |\tilde{1}\rangle &= \sin\frac{\pi}{8}|0\rangle + \cos\frac{\pi}{8}|1\rangle. \end{aligned} \tag{13}$$

Armed with this observable, Alice has a certain probability $\mathfrak{p}_{ch}(a|b)$ of obtaining the result a whenever state $|b\rangle$ was sent. Concretely, in an r-noise channel, we have:

$$\begin{aligned} \mathfrak{p}_{ch}(0|0) &= \frac{r}{2} + \frac{1}{4}(2 + \sqrt{2} - (2+\sqrt{2})r) \\ \mathfrak{p}_{ch}(1|0) &= \frac{r}{2} + \frac{1}{4}(2 - \sqrt{2} - (2-\sqrt{2})r) \\ \mathfrak{p}_{ch}(0|+) &= \frac{1}{4}(2 - \sqrt{2} + r\sqrt{2}) \\ \mathfrak{p}_{ch}(1|+) &= \frac{1}{4}(2 + \sqrt{2} - r\sqrt{2}). \end{aligned} \tag{14}$$

Note that the Breidbart basis, optimal for distinguishing the states $|0\rangle$ and $|+\rangle$, is also optimal for distinguishing the mixed states obtained as a result of the action of a depolarizing channel (due to the symmetry of white noise) as it saturates the Helstrom bound.

Once Alice has inferred the state of every photon, she can send her results back to Bob in hopes of passing his test. Such a strategy has been considered and analysed in previous work [10]. However, she can do better. After all, Alice knows that her inference is not completely accurate, and so sending her inference on what the states sent by Bob may be is definitely not the best she can do. In fact, even if Alice were able to perfectly determine the state of Bob's particles, she wouldn't send these results back to Bob, but rather results that would emulate the statistics of an honest Alice (such as (2) or (5)). She can thus proceed as follows: every time she obtains a 0 with the cheating observable, she decides to either change the result, with a certain probability $p_{0\to1}$, or to keep it, with probability $(1 - p_{0\to1})$. Analogously, every time Alice obtains a 1, she decides either to change it, with probability $p_{1\to0}$, or to keep it, with probability $(1 - p_{1\to0})$. After such a procedure, denoting by $p_{ch}(a|b)$ the probability that Alice sends the value a back to Bob whenever

he had sent $|b\rangle$, we have:

$$\begin{aligned} p_{ch}(0|0) &= \mathfrak{p}_{ch}(0|0)\cdot(1-p_{0\to1}) + \mathfrak{p}_{ch}(1|0)\cdot p_{1\to0} \\ p_{ch}(1|0) &= \mathfrak{p}_{ch}(1|0)\cdot(1-p_{1\to0}) + \mathfrak{p}_{ch}(0|0)\cdot p_{0\to1} \\ p_{ch}(0|+) &= \mathfrak{p}_{ch}(0|1)\cdot(1-p_{0\to1}) + \mathfrak{p}_{ch}(1|0)\cdot p_{1\to0} \\ p_{ch}(1|+) &= \mathfrak{p}_{ch}(1|1)\cdot(1-p_{1\to0}) + \mathfrak{p}_{ch}(0|1)\cdot p_{0\to1}. \end{aligned} \tag{15}$$

The next step is rather straightforward: determine the values of $p_{0\to1}$ and $p_{1\to0}$ that most increase Alice's chances of successfully cheating. Naturally, this depends on Bob's criterion, and on whether Alice is trying to trick Bob into believing she committed to 0, or trying to trick him into believing she committed to 1. Without loss of generality, assume Alice is trying to forge a commitment to 0, and that Bob is using the goodness-of-fit test described in Section 2, that is he accepts Alice's commitment to 0 if the number of 0's revealed by Alice falls within the interval $[\mu_0 - 3\sigma_0, \mu_0 + 3\sigma_0]$ whenever $|0\rangle$ was sent (see (9)), and the number of 1's within the interval $[\mu_+ - 3\sigma_+, \mu_+ + 3\sigma_+]$ whenever $|+\rangle$ was sent (see (11)). The probability that a cheating Alice passes this test is then given by

$$p_{ch} = \left(\sum_{k=\mu_0-3\sigma_0}^{\mu_0+3\sigma_0} \binom{N_0}{k} p_{ch}^{k}(0|0)\cdot p_{ch}^{N_0-k}(1|0)\right)\cdot\left(\sum_{k=\mu_+-3\sigma_+}^{\mu_++3\sigma_+} \binom{N_+}{k} p_{ch}^{k}(0|+)\cdot p_{ch}^{N_1-k}(1|+)\right) \tag{16}$$

where $N_0 = N_+ = N$ are the number of $|0\rangle$ states and the number of $|+\rangle$ states sent by Bob, respectively.

Although finding an expression for the maximum of such a function is non-trivial, it is quite simple to do so numerically. In Fig. 2 we present a few plots of p_{ch}, for different values of the noise parameter r, and for $M = N_0 + N_+ = 2N = 100$.

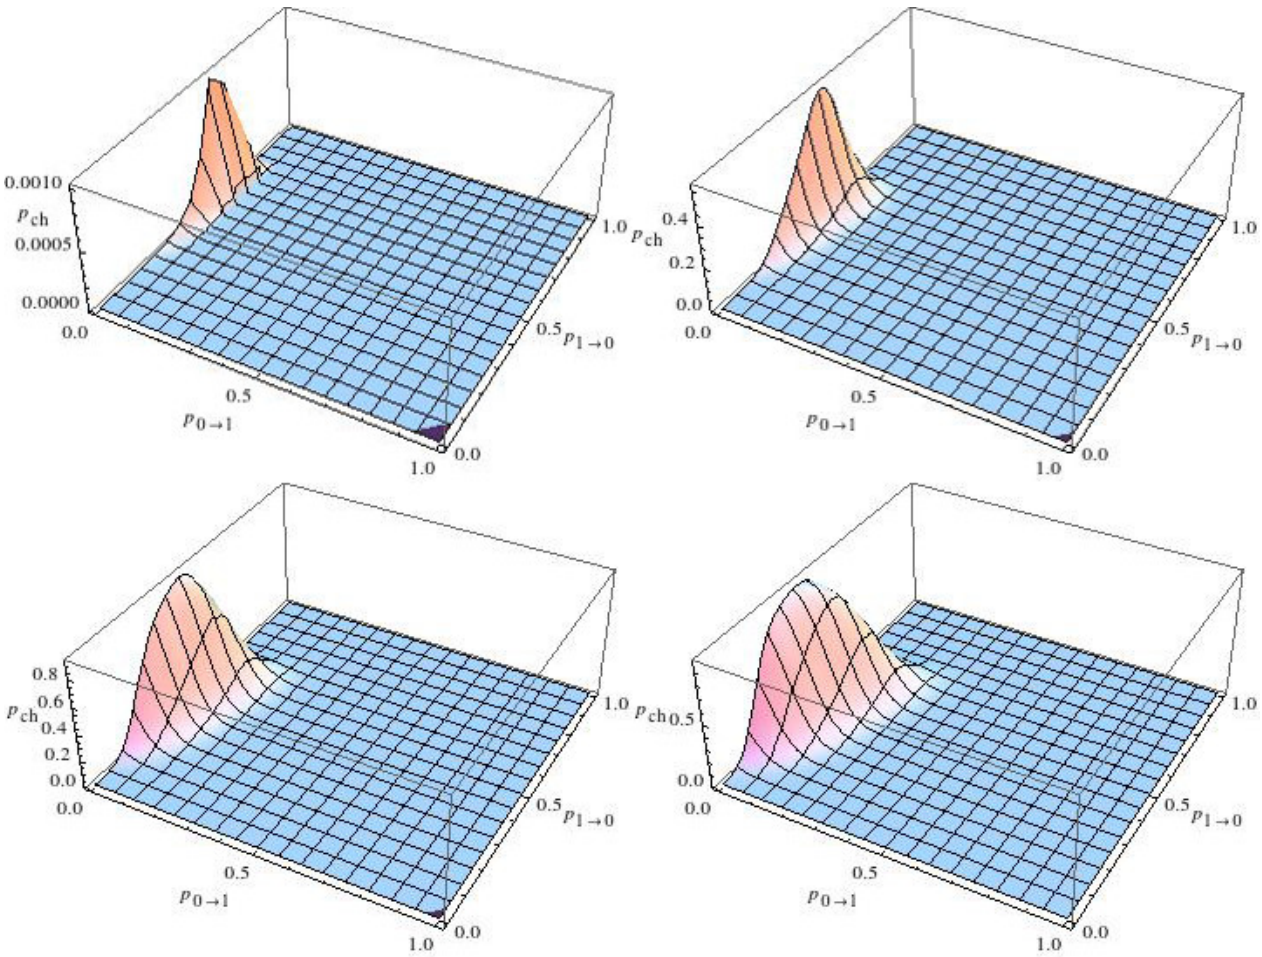

Figure 2: Probability p_{ch} of a cheating Alice with $M = N_0 + N_+ = 2N = 100$ successful measurements passing Bob's test of commitment to 0, for different values of the noise parameter r (0 - top left, 0.08 - top right, 0.16 - bottom left, 0.24 - bottom right).

As can be seen, to obtain an optimal cheating strategy in the case of a noiseless channel, a cheating Alice should always choose $p_{0\to1}$ to be 0. For higher values of noise Alice should choose $p_{0\to1}$ to be slightly higher than 0. For $M = 100$ the threshold value for r, for which Alice should choose $p_{0\to1} > 0$, is above 0.52, and it decreases as M grows. As for $p_{1\to0}$, the optimal value lies in the interval $(0.4, 0.5)$, but depends on the noise parameter r. Given the optimal choice of the parameters $p_{0\to1}$ and $p_{1\to0}$, the maximal probability to cheat, $p_{ch}^{\max}$, is presented in Fig. 3, as a function of $M = N_0 + N_+ = 2N$ and r.

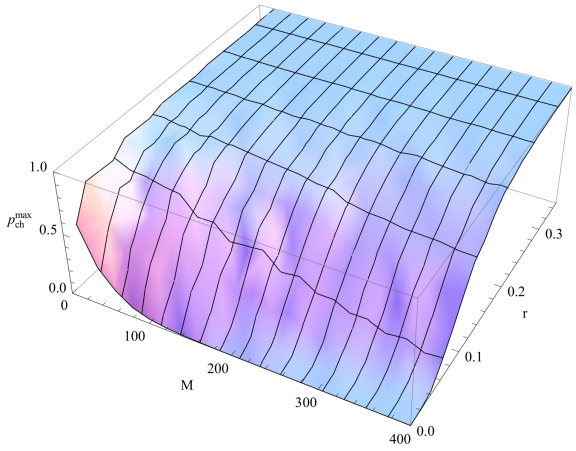

Figure 3: Probability $p_{ch}^{\max}$ of a cheating Alice as a function of the number $M = N_0 + N_+ = 2N$ of successful measurements passing Bob's test of commitment to 0, and the noise parameter r.

3.2 The four-state protocol

It is quite interesting to compare the 2-state cheating strategy above with the analogous cheating strategy for the 4-state variant of the protocol. This variant, initially proposed in [4], goes as follows:

1. Alice and Bob agree on the two observables

$$\begin{aligned} \hat{C}_0 &= 0 \cdot |0\rangle\langle 0| + 1 \cdot |1\rangle\langle 1| \\ \hat{C}_1 &= 0 \cdot |-\rangle\langle -| + 1 \cdot |+\rangle\langle +|. \end{aligned} \tag{17}$$

2. Bob randomly chooses one of the four states, $|0\rangle$, $|1\rangle$, $|+\rangle$ of $|-\rangle$, and then prepares and sends a particle in that state to Alice. He repeats this step a number M of times of his choice.

3. Alice measures either $\hat{C}_0$, for committing to 0, on all the particles she receives; or she measures $\hat{C}_1$, for committing to 1, on all the particles she receives.

4. Whenever Alice wants to reveal her commitment, she sends her commitment choice along with the results of her measurements back to Bob.

5. Bob either accepts or discards Alice's commitment, based on some statistical criterion of his choosing.

Instead of two tests, Alice now has to pass four distinct tests: one for whenever each different state was sent. The four probability distributions would now be given by Equations (5), (6), and two additional but analogous ones for whenever states $|1\rangle$ and $|-\rangle$ are sent. Again, in Fig. 4 we plot the probability that a cheating Alice, measuring all photons in the Breidbart basis (the same as for the 2-state protocol), passes all four tests of commitment to 0, as a function of the probability to change her mind about 0's ($p_{0\to 1}$) and about 1's ($p_{1\to 0}$), for $N = N_{0/1} = N_{+/-} = 50$.

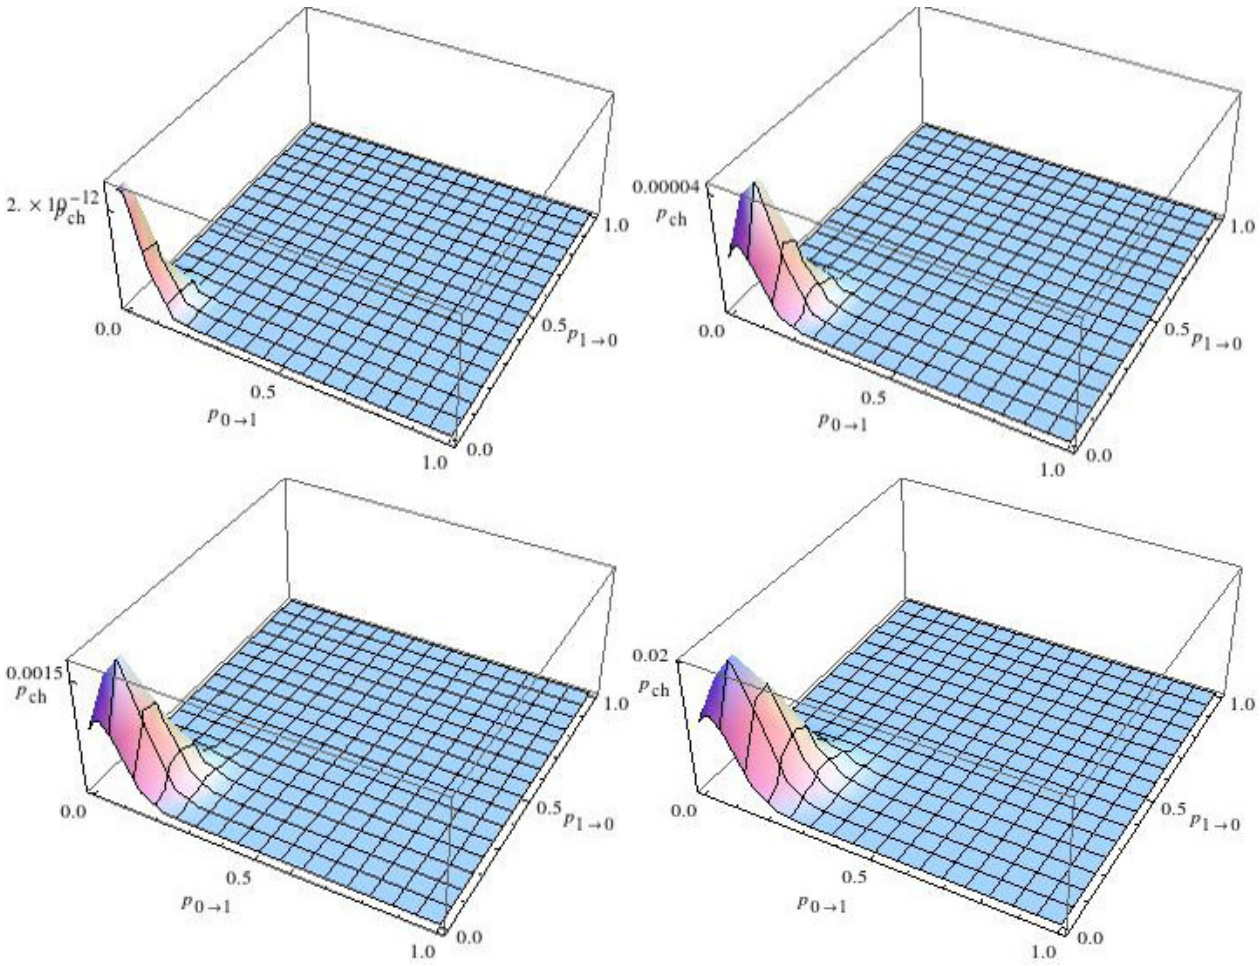

Figure 4: Probability of a cheating Alice with $M = N_0 + N_1 + N_+ + N_- = 200$ successful measurements passing Bob's test of commitment to 0, for different values of the noise parameter r (0 - top left, 0.08 - top right, 0.16 - bottom left, 0.24 - bottom right).

As can be seen, it is always advantageous for a cheating Alice to change her mind about both 0's and 1's. Finally, in Fig. 5 we show the maximal probability to cheat, $p_{ch}^{\max}$, as a function of $M = N_0 + N_1 + N_+ + N_- = 4N$ and r. Comparing with the 2-state case, it is clear that in terms of the resources (the total number of measurement outcomes, $M = 2N$ for two-state, and $M = 4N'$ for four-state protocols, with $N' = N/2$) the four-state protocol is advantageous.

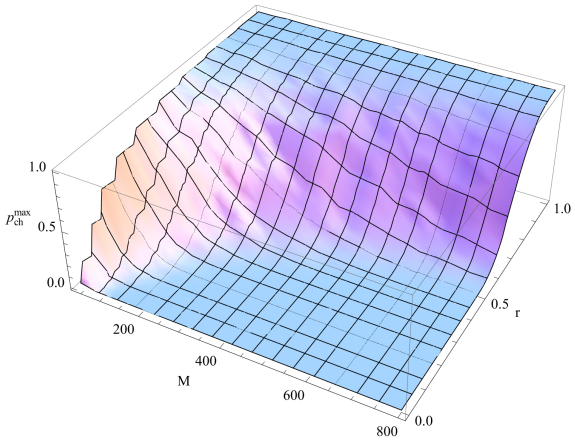

Figure 5: Probability $p_{ch}^{\max}$ of a cheating Alice as a function of the number $M = 4N$ of successful measurements passing Bob's test of commitment to 0, and the noise parameter r.

4 Faking-distance attack

In any experimental implementation of the above protocols, there will be inevitable losses during transmission, just as there is inevitable noise. These losses, just like noise, need to be taken into account by Bob when receiving the final data from Alice. Indeed, if Alice is close by, then Bob expects her to receive most of the emitted particles, say, photons, and thus expects a measurement result for almost each photon sent. On the other hand, if Alice is far away, then several photons will be lost during transmission, and thus Bob will only be expecting a small number of measurement results when compared to the number of photons sent. Curiously enough, this mundane phenomenon is advantageous for a cheating Alice.

Suppose Alice is much closer to Bob than she claimed to be, meaning Bob is expecting that only a fraction f of the particles sent will end up producing a measurement result. In reality, Alice is able to generate measurement results for (almost) all emitted photons. So if Alice has a way to produce statistical data in perfect agreement with those of two honest parties committing to different values for a fraction f of all particles, she can trick Bob into thinking her forged statistics are the legitimate statistics obtained by an Alice far away, who had access to only f times the particles sent.

Alice can achieve this very easily by simply measuring $\hat{C}_0$ on half of the received particles, and $\hat{C}_1$ on the other half. When asked to give proof of her commitment, Alice simply sends the half of the measurement results that best suits her, thus emulating an honest Alice who would have had access to only 50% of the particles sent. This strategy is stronger than the equally common technique of unambiguous state discrimination, where Alice can learn the true state of around 29% of the particles she receives.

A cheating Alice, making her measurements right outside Bob's lab, while pretending to be far away, can thus cheat if Bob is expecting that only 50% or less of the particles sent end up producing a measurement result (by an honest Alice). There is thus a maximal distance up to which the protocol is secure: the distance beyond which an honest Alice is only able to measure 50% or less of the particles sent. To estimate this distance, let N_P be the total number of pulses produced by Bob. In each pulse, a photon is emitted following a certain probability distribution with mean μ. The total number of emitted photons is then $N_E = \mu N_P$. The number of photons received by an honest Alice is thus $N_R = 10^{-\alpha L/10} N_E$, where α is the attenuation coefficient and L is the length of the optical fibre. Finally, due to a finite detector efficiency η, the number of

measurement outcomes of an honest Alice is $N_M = \eta N_R = \eta\mu 10^{-\alpha L/10} N_P$, while for a cheating Alice it is $N_{cM} = \eta N_E = \eta\mu N_P$. A cheating Alice, performing measurements right outside Bob's lab ($L = 0$), produces statistics that are compatible with her future choice of commitment on 50% of the measurement outcomes. Note that we are assuming that a cheating Alice has the same type of detector as an honest Alice, since we are dealing with a practical case scenario. Thus, the maximal distance $L_{\max}$ for which the protocol is still secure is given by $N_M = N_{cM}/2$, leading to

$$L_{\max} = \frac{10}{\alpha} \log_{10}(2). \tag{18}$$

For a typical value of $\alpha = 0.2$ (as used in, for instance, [2]), this amounts to a maximal safe distance of 15km.

This attack shares quite some similarity with the famous photon-number splitting attack on the BB84 protocol, and which can be patched using decoy states, as suggested by Hwang in [6]. Unfortunately, this remedy is not applicable in our case, since the evildoer (Eve) and the honest agent (Alice) are one and the same. We are thus left with only a maximum safe distance.

In the case of real-life noisy scenario, it is natural to assume that the "distant" Alice obtains imperfect results due to the "distant" noise factor r_d, while the cheating "nearby" Alice has noise given by $r_n < r_d$. A cheating Alice has therefore better statistics than an honest one, and can use her "good" results on 50% of the cases to infer, say $(50 + \delta)$% of the "less quality" results obtained by an honest Alice: by randomly choosing δ% of the "missing" results, a cheating Alice spoils the overall quality of her statistics, to match those of an honest Alice, which additionally decreases $L_{\max}$. In other words, we have

$$N_M = \frac{N_{cM}}{2} + \Delta N, \tag{19}$$

where $\Delta N = \delta \cdot N$ (consequently, $\delta = 1/2 - 10^{-\alpha L/10}$). The statistics $p_{c;r_d}(*|*)$, with $c \in \{0, 1\}$, of an honest Alice are given by Equations (5) and (6), with $r = r_d$. The statistics of a cheating Alice are given by a weighted sum of the "good" results, measured on $N_{cM}/2$ photons, and random guesses, presented for ΔN photons:

$$p^{ch}_{c;r_n}(*|*) = \frac{\frac{N_{cM}}{2}}{\frac{N_{cM}}{2} + \Delta N} p_{c;r_n}(*|*) + \frac{\Delta N}{\frac{N_{cM}}{2} + \Delta N} \frac{1}{2}. \tag{20}$$

Equating, say $p^{ch}_{0;r_n}(0|0)$ with $p_{0;r_d}(0|0)$, and using (19), one obtains the expression for the maximum safe distance to be

$$L_{\max} = \frac{10}{\alpha} [\log_{10}(2) + \log_{10}(1 - \frac{r_d - r_n}{1 - r_d})]. \tag{21}$$

Note that the above expression (21) is defined only for $(r_d - r_n)/(1 - r_d) < 1$. If $r_d \geq (r_n + 1)/2$, there exist no $L_{\max}$, i.e., a dishonest Alice can always fake her results successfully, which comes as no surprise – for such high noise $r_d > 1/2$ the two commitments become indistinguishable and the protocol is not viable. The above analysis is applicable to both two- and four-state protocols.

5 Multi-photon sources

There is still one detail that we've up until now discarded but which plagues any implementation that may nowadays be done: the nonexistence of perfect single-photon sources. Indeed, up until this point we've never considered the possibility that Alice may find with certainty in which state are certain particles sent by Bob. It turns out, however, that any device used to prepare photons

in a specific state has a non-negligible probability of creating a pair (or even triple, quadruple, ...) of photons in the same state instead. In a realistic scenario, one must consider a photon source whose number of photons emitted per pulse follows a Poisson distribution with a certain parameter μ (average number of photons per pulse), characteristic of the source. If we denote by $P(n,\mu)$ the probability to obtain n photons in a given pulse, we have

$$P(n,\mu) = \frac{\mu^n}{n!}e^{-\mu}. \tag{22}$$

Alas, as soon as Bob sends more than one photon in a single pulse, which happens with probability $(1 - e^{-\mu} - \mu e^{-\mu})$, then Alice may split up the photons and measure the two observables, $\hat{C}_0$ and $\hat{C}_1$, on two of them separately, thus learning the photon's true state (provided Bob follows the protocol) and producing a statistics compatible with either a commitment to 0 or a commitment to 1 upon request. If Alice receives but one photon, she reverts back to the strategy described in Section 3. We must then once again reevaluate the sets of cheating probabilities in order to obtain $p_{CH}(i|j)$ - the probability that a cheating Alice, using every single tool at her disposal, will send the value i back to Bob, every time he sent state $|j\rangle$. When trying to forge a commitment to 0, the corresponding set of conditional probabilities for a two-state case is (the additional conditional probabilities for a four-state case are obtained analogously)

$$\begin{aligned}
p_{CH}(0|0) &= \frac{1}{1-e^{-\mu}}[\mu e^{-\mu} p_{ch}(0|0) + (1 - e^{-\mu} - \mu e^{-\mu}) p_0(0|0)] \\
p_{CH}(1|0) &= \frac{1}{1-e^{-\mu}}[\mu e^{-\mu} p_{ch}(1|0) + (1 - e^{-\mu} - \mu e^{-\mu}) p_0(1|0)] \\
p_{CH}(0|+) &= \frac{1}{1-e^{-\mu}}[\mu e^{-\mu} p_{ch}(0|+) + (1 - e^{-\mu} - \mu e^{-\mu}) p_0(0|+)] \\
p_{CH}(1|+) &= \frac{1}{1-e^{-\mu}}[\mu e^{-\mu} p_{ch}(1|+) + (1 - e^{-\mu} - \mu e^{-\mu}) p_0(1|+)],
\end{aligned} \tag{23}$$

and analogously for the case of a commitment to 1, by replacing $p_0(*|*)$ with $p_1(*|*)$.

In more detail, take for instance the very first value: $p_{CH}(0|0)$. This value only has meaning if Bob actually sent something, which happens with probability $(1 - e^{-\mu})$, hence the factor $\frac{1}{1-e^{-\mu}}$. Given the case, two mutually incompatible results may occur. Either Bob sends a single particle, which happens with probability $\mu e^{-\mu}$, and Alice measures it using her single-photon strategy from Section 3 (hence $p_{ch}(0|0)$), or Bob sends two or more particles, which happens with probability $(1 - e^{-\mu} - \mu e^{-\mu})$, and Alice perfectly emulates the probability distribution of an honest party (hence $p_0(0|0)$). All other expressions follow an analogous reasoning.

Note that that the conditional probabilities $p_{ch}(*|*)$, corresponding to the strategy from Section 3, are those which maximise the overall cheating probability p_{ch}^{max} with respect to the case of the *multi-photon* sources. Thus, the optimal values of $p_{0\to1}$ and $p_{1\to0}$ are, in general, different for the strategies that maximise the cheating probability for the case of the *single-photon* source, and the case of the *multi-photon* source. In Tables 1 and 2, we compare the mentioned optimal values for both two- and four-state protocols, respectively, for $M = 100, 200, 300$ and 400, with $r = 0.1$ and, for the multi-photon case, $\mu = 0.2$ (the value used in the experimental realisation of a two-state protocol [2]).

The cheating strategy considered here is beyond the capabilities of technology, both in the present and in any foreseeable future. To perform this kind of attack, a cheating Alice needs far more than just a non-demolition measurement. She needs a device capable of detecting whether or not the photon source emitted a single photon, and needs to change her measurement apparatus accordingly. If a single photon was sent, she measures it using the optimal cheating observable.

M	Single-photon	Multi-photon
100	(0.0, 0.489663)	(0.0, 0.492572)
200	(0.0, 0.490563)	(0.0, 0.494314)
300	(0.0, 0.479936)	(0.0, 0.483053)
400	(0.0, 0.47544)	(0.0, 0.478355)

Table 1: The two-state protocol's optimal values $(p_{0\to1}, p_{1\to0})$ for single-photon (left column) and multi-photon (right column) sources, for $M = 100, 200, 300$ and 400, with $r = 0.1$ and, for the multi-photon case, $\mu = 0.2$.

M	Single-photon	Multi-photon
200	(0.0658591, 0.0658591)	(0.0470355, 0.0470355)
400	(0.0793028, 0.0793028)	(0.0592754, 0.0592754)
600	(0.0939039, 0.0939039)	(0.0743531, 0.0743531)
800	(0.101166, 0.101166)	(0.0818707, 0.0818707)

Table 2: The four-state protocol's optimal values $(p_{0\to1}, p_{1\to0})$ for single-photon (left column) and multi-photon (right column) sources, for $M = 100, 200, 300$ and 400, with $r = 0.1$ and, for the multi-photon case, $\mu = 0.2$.

Otherwise, she splits two of the incoming photons into two different measurement apparatuses (measuring $\hat{C}_0$ and $\hat{C}_1$ respectively). Such technology is far beyond our current capabilities. More realistic is to perform the beam-splitter strategy to all of the pulses. This way, a cheating Alice would on average measure $\hat{C}_0$ on half of the single-photon pulses ($\frac{1}{2}\mu e^{-\mu} N_P$), and $\hat{C}_1$ on the other half, while for the rest of the multi-photon pulses she will have the results of both the observables. Thus, for the photons on which the wrong observable was measured, a cheating Alice has to add random results. When committing to 0, her statistics for a two-state case (the additional conditional probabilities for a four-state case are obtained analogously) are given by

$$p_{BS}(*|*) = \left(1 - \frac{1}{2}\frac{\mu e^{-\mu}}{1-e^{-\mu}}\right) p_0(*|*) + \left(\frac{1}{2}\frac{\mu e^{-\mu}}{1-e^{-\mu}}\right) \cdot \frac{1}{2}, \tag{24}$$

and analogously for a commitment to 1, by replacing $p_0(*|*)$ with $p_1(*|*)$.

Naturally, the higher μ is, the closer the above above two distributions given by (23) and (24) are to those obtained by an honest Alice committing to 0 or committing to 1, respectively. Thus, we only consider weak sources with low average photon number $\mu \in \{0, 1\}$ (indeed, in a recent experimental realisation of a two-state protocol [2], the value $\mu \approx 0.2$ was used). In Figures 6, 7, 8 and 9 are the plots of the probabilities to cheat for both types of the strategies (less realistic (23) and more realistic (24)), and both versions of the protocol (two- and four-state), as functions of the average number of photons per pulse μ and the noise parameter r, for $M = 100, 200, 300$ and 400, respectively.

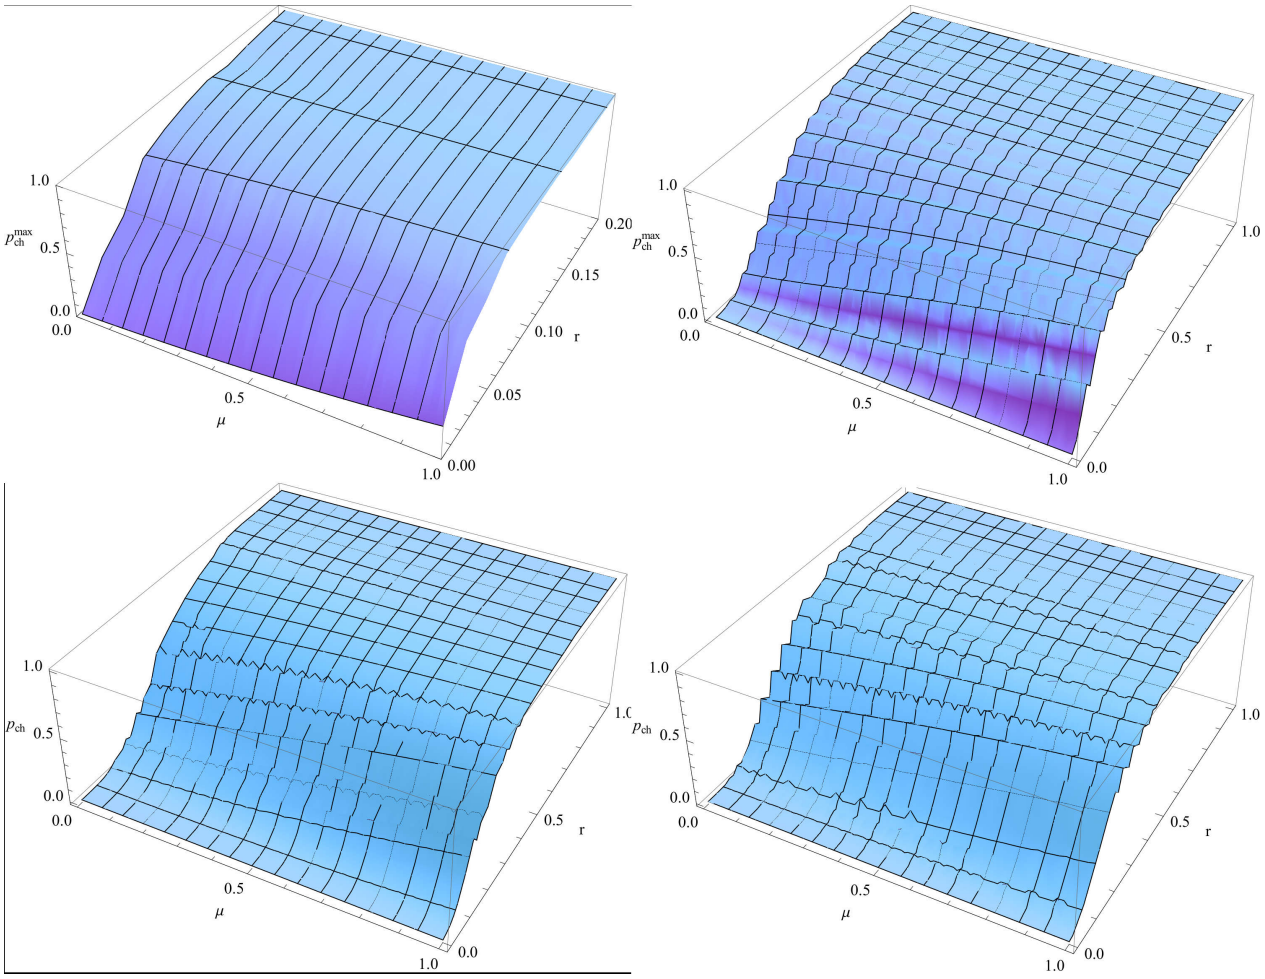

Figure 6: Probabilities of a cheating Alice with $M = 100$ successful measurements passing Bob's test of commitment to 0, as functions of the average number of photons per pulse $\mu \in \{0, 1\}$ and the noise parameter r, for: the two-state protocol with the less realistic (23) strategy (top left); the two-state protocol with more realistic (24) strategy (bottom left); the four-state protocol with the less realistic (23) strategy (top right); the four-state protocol with more realistic (24) strategy (bottom right).

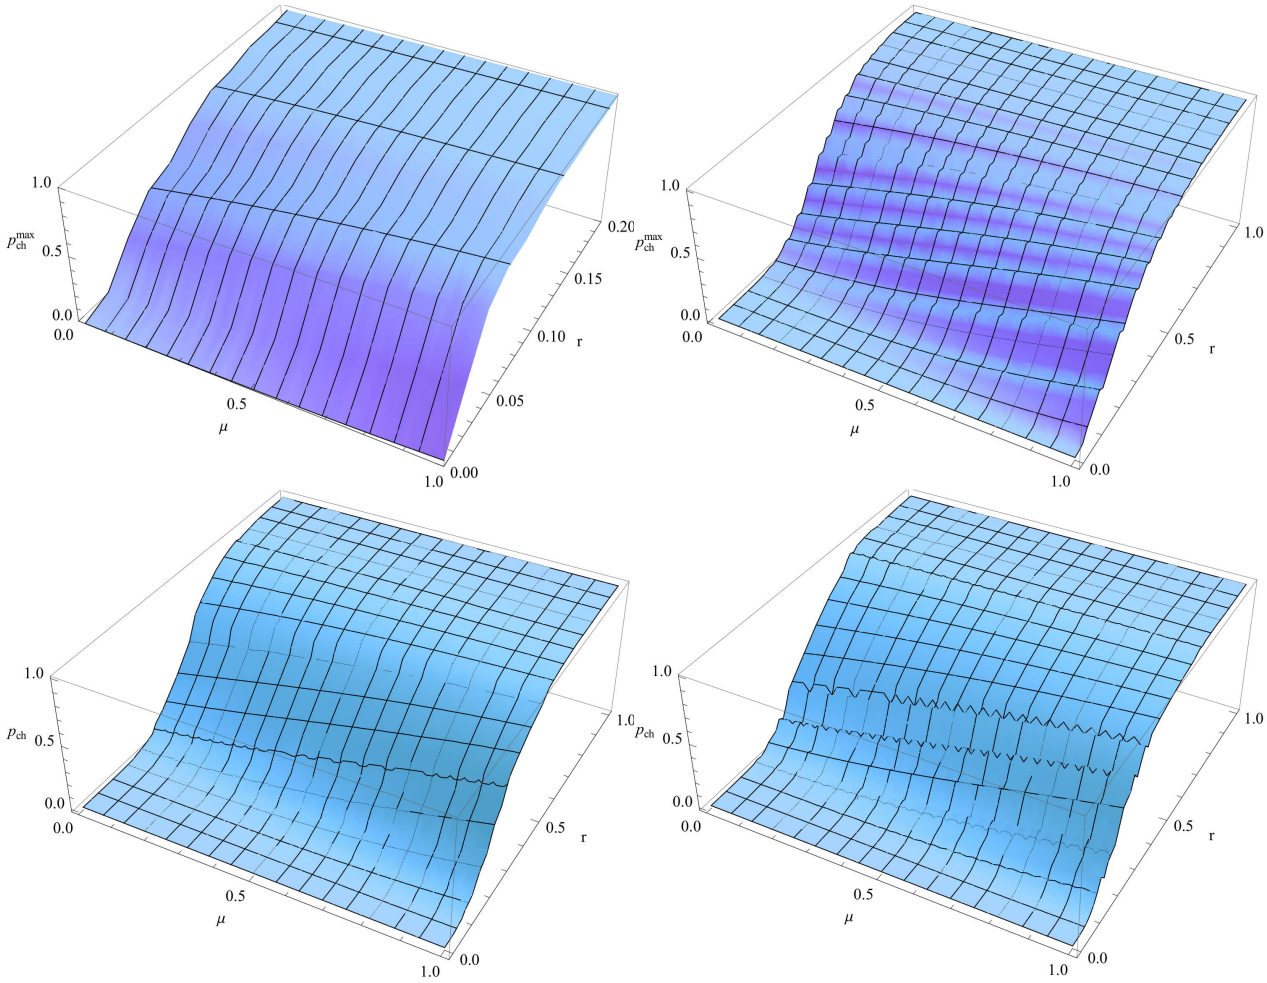

Figure 7: Probabilities of a cheating Alice with $M = 200$ successful measurements passing Bob's test of commitment to 0, as functions of the average number of photons per pulse $\mu \in \{0, 1\}$ and the noise parameter r, for: the two-state protocol with the less realistic (23) strategy (top left); the two-state protocol with more realistic (24) strategy (bottom left); the four-state protocol with the less realistic (23) strategy (top right); the four-state protocol with more realistic (24) strategy (bottom right).

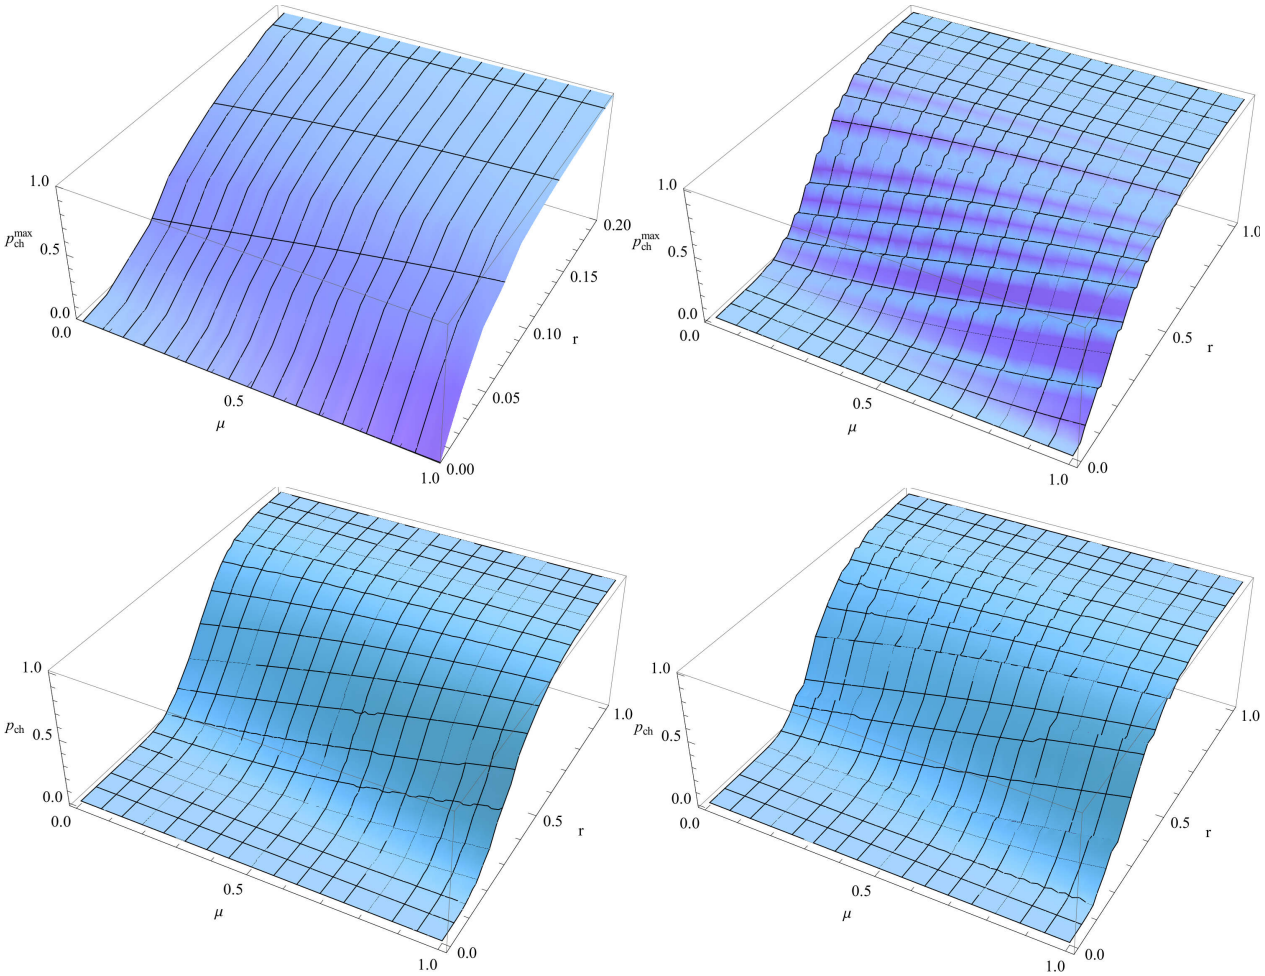

Figure 8: Probabilities of a cheating Alice with $M = 300$ successful measurements passing Bob's test of commitment to 0, as functions of the average number of photons per pulse $\mu \in \{0, 1\}$ and the noise parameter r, for: the two-state protocol with the less realistic (23) strategy (top left); the two-state protocol with more realistic (24) strategy (bottom left); the four-state protocol with the less realistic (23) strategy (top right); the four-state protocol with more realistic (24) strategy (bottom right).

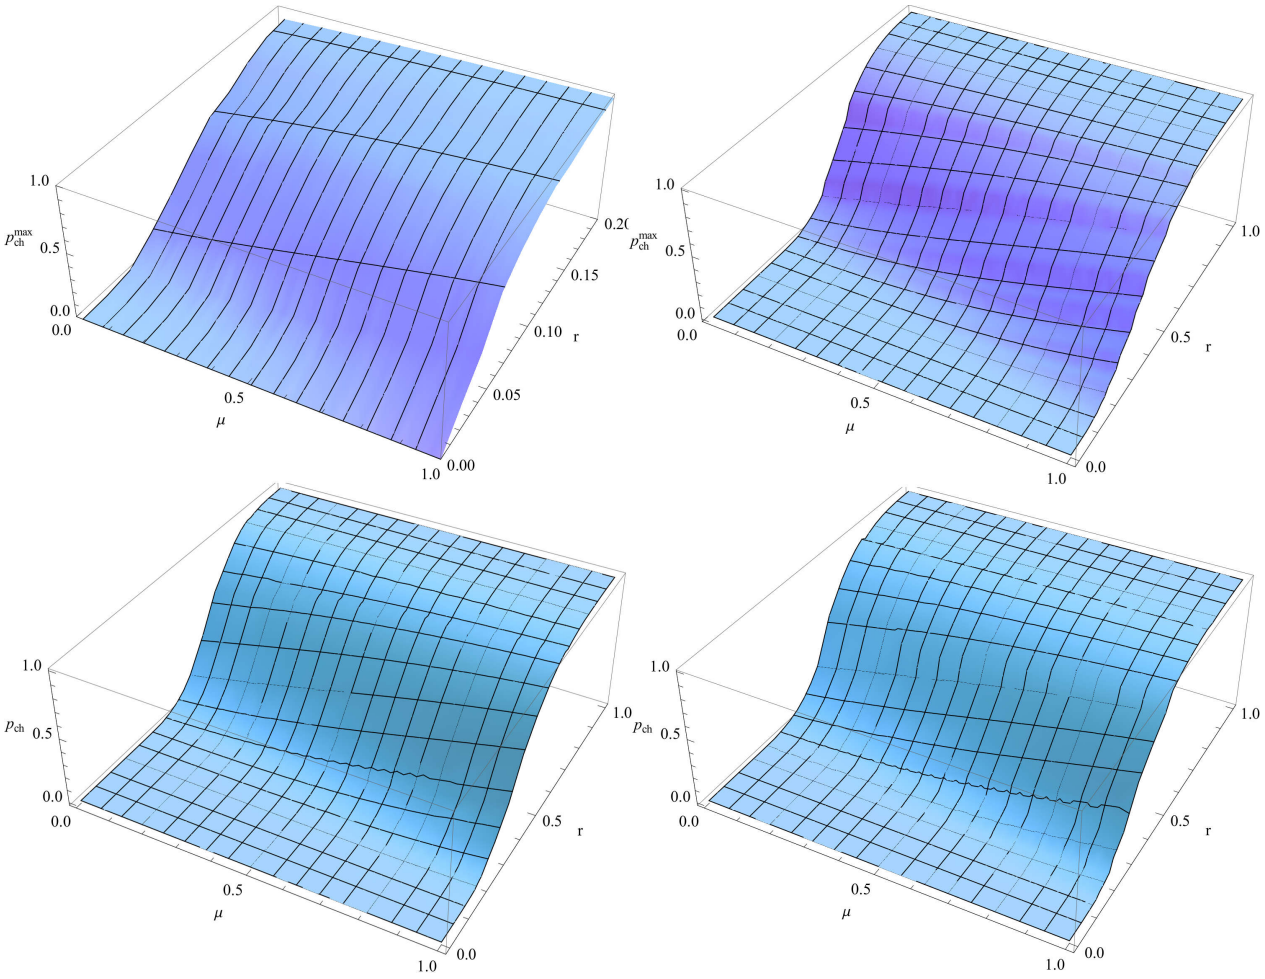

Figure 9: Probabilities of a cheating Alice with $M = 400$ successful measurements passing Bob's test of commitment to 0, as functions of the average number of photons per pulse $\mu \in \{0, 1\}$ and the noise parameter r, for: the two-state protocol with the less realistic (23) strategy (top left); the two-state protocol with more realistic (24) strategy (bottom left); the four-state protocol with the less realistic (23) strategy (top right); the four-state protocol with more realistic (24) strategy (bottom right).

6 Conclusions

We have analysed the cheating strategies for practical two-state and four-state bit-commitment protocols presented in [10] and [4], respectively. We showed that by introducing the "post-processing" of the raw experimental data, one can improve the naïve cheating strategy based on the measurement in the Breidbart basis, originally studied in [10]. We studied the two cases of perfect single-photon and realistic multi-photon sources, showing that the "post-processing" optimal parameters $p_{0\to1}$ and $p_{1\to0}$ are in general different for the two cases. We have also analysed the more realistic multi-photon source beam-splitter strategy that does not rely on (single-photon) non-demolition measurements. In all of the mentioned cases, the four-state protocol shows clear advantage over the two-state version, with respect to the resources needed to achieve the same level of security. Finally, we analysed an alternative strategy, the *faking-distance attack*, based on the forged geographical location of the committing party (Alice), showing that for today's typical equipment the protocol

is secure up to a distance of 15km between the parties, both for two- and four-state protocols. Our approach is, following the analysis presented in [10], straightforward to apply for the cases of noisy and/or bounded memories.

Acknowledgments

The authors acknowledge bilateral scientific co-operation between Portugal and Serbia through the project "Noise and measurement errors in multi-party quantum security protocols", no. 451-03-01765/2014-09/04 supported by the Foundation for Science and Technology (FCT), Portugal, and the Ministry of Education, Science and Technological Development of the Republic of Serbia. RL thanks the support of FCT through the PhD Grant SFRH/BD/79571/2011. NP acknowledges the support of SQIG – Security and Quantum Information Group and IT project QbigD funded by FCT PEst-OE/EEI/LA0008/2013 and UID/EEA/50008/2013. DA, DP and SP acknowledge the support of the Ministry of Education and Science of the Republic of Serbia, Contract No. 171006, No. 171017, No. 171020, No. 171038 and No. 45016.

References

[1] Á. J. Almeida, R. Loura, N. Paunković, N. A. Silva, N. J. Muga, P. Mateus, P. S. André, and A. N. Pinto. A brief review on quantum bit commitment. In *Proc. SPIE*, pages 92861C–92861C–8, 2014.

[2] Á. J. Almeida, A. D. Stojanovic, N. Paunković, R. Loura, N. J. Muga, N. A. Silva, P. Mateus, P. S. André, and A. N. Pinto. Implementation of a two-state quantum bit commitment protocol in optical fibers. *Journal of Optics*, 18(1):015202, 2016.

[3] M. Blum. Coin flipping by telephone. In *Proc. of Crypto*, pages 11–15, 1981.

[4] A. Danan and L. Vaidman. Practical quantum bit commitment protocol. *Quantum Information Processing*, 11(3):769–775, 2012.

[5] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in np have zero-knowledge proof systems. *J. ACM*, 38(3):690–728, July 1991.

[6] Won-Young Hwang. Quantum key distribution with high loss: Toward global secure communication. *Phys. Rev. Lett.*, 91:057901, Aug 2003.

[7] R. Koenig, S. Wehner, and J. Wullschleger. Unconditional security from noisy quantum storage. *IEEE Transactions on Information Theory*, 58(3):1962–1984, 2012.

[8] Yehuda Lindell and Benny Pinkas. *An Efficient Protocol for Secure Two-Party Computation in the Presence of Malicious Adversaries*, pages 52–78. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[9] H. K. Lo and H. F. Chau. Is quantum bit commitment really possible? *Phys. Rev. Lett.*, 78:3410–3413, 1997.

[10] R. Loura, Á. J. Almeida, P. S. André, A. N. Pinto, P. Mateus, and N. Paunković. Noise and measurement errors in a practical two-state quantum bit commitment protocol. *Phys. Rev. A*, 89:052336, May 2014.

[11] D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Physycal Review Letters*, page 3414, 1997.

[12] N. Ng, S. Joshi, C. Ming, C. Kurtsiefer, and S. Wehner. Experimental implementation of bit commitment in the noisy-storage model. *Nature Communications*, 3:1326–, 2012.

[13] C. Schaffner, B. M. Terhal, and S. Wehner. Robust cryptography in the noisy-quantum-storage model. *Quantum Information & Computation*, 9(11):963–996, 2011.

[14] A. Shamir, R. L. Rivest, and L. M. Adleman. Mental poker. In D. A. Klarner, editor, *The Mathematical Gardner*, pages 37–43. Springer US, 1981.

[15] A. Souto, P. Mateus, P. Adão, and N. Paunković. Bit-string oblivious transfer based on quantum state computational distinguishability. *Phys. Rev. A*, 91:042306, 2015.

[16] S. Wehner, C. Schaffner, and B. M. Terhal. Cryptography from noisy storage. *Phys. Rev. Lett.*, 100:220502, 2008.